

Internet Draft
Expires: August 12, 2012

H. Kaplan (ed.)
Acme Packet
K. Hedayat
EXFO
N. Venna
Saperix
P. Jones
Cisco Systems, Inc.
A. Roychowdhury
Hughes Systique Corp.
C. SivaChelvan
Cisco Systems, Inc.
N. Stratton
BlinkMind, Inc.
March 26, 2012

**An Extension to the Session Description Protocol (SDP)
and Real-time Transport Protocol (RTP) for Media Loopback
draft-ietf-mmusic-media-loopback-18**

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 12, 2012.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Simplified BSD License.

Abstract

The wide deployment of Voice over IP (VoIP), Text and Video over IP services has introduced new challenges in managing and maintaining real-time voice/text/video quality, reliability, and overall performance. In particular, media delivery is an area that needs attention. One method of meeting these challenges is monitoring the media delivery performance by looping media back to the transmitter. This is typically referred to as "active monitoring" of services. Media loopback is especially popular in ensuring the quality of transport to the edge of a given VoIP, Real-time Text or Video over IP service. Today in networks that deliver real-time media, short of running 'ping' and 'traceroute' to the edge, administrators are left without the necessary tools to actively monitor, manage, and diagnose quality issues with their service. The extension defined herein adds new SDP media attributes, which enable establishment of media sessions where the media is looped back to the transmitter. Such media sessions will serve as monitoring and troubleshooting tools by providing the means for measurement of more advanced VoIP, Real-time Text and Video over IP performance metrics.

Table of Contents

1. Introduction.....	3
1.1 Use Cases Supported.....	4
2. Terminology.....	6
3. Overview of Operation.....	6
3.1 SDP Offerer Behavior.....	6
3.2 SDP Answerer Behavior.....	6
4. New SDP Attributes.....	7
4.1 Loopback Type Attribute.....	7

4.2 Loopback Role Attributes: loopback-source and loopback-mirror.....	8
5. Rules for Generating the SDP offer/answer.....	9
5.1 Generating the SDP Offer for Loopback Session.....	9
5.2 Generating the SDP Answer for Loopback Session.....	10
5.3 Offerer Processing of the SDP Answer.....	11
5.4 Modifying the Session.....	12
5.5 Establishing Sessions Between Entities Behind NAT.....	12
6. RTP Requirements.....	12
7. Payload formats for Packet loopback.....	13
7.1 Encapsulated Payload format.....	13
7.2 Direct loopback RTP payload format.....	16
8. RTCP Requirements.....	17
9. Congestion Control.....	17
10. Examples.....	18
10.1 Offer for specific media loopback type.....	18
10.2 Offer for choice of media loopback type.....	18
10.3 Answerer rejecting loopback media.....	19
11. Security Considerations.....	20
12. Implementation Considerations.....	21
13. IANA Considerations.....	21
13.1 SDP Attributes.....	21
13.2 MIME Types.....	22
14. Acknowledgements.....	31
15. Normative References.....	31
16. Informative References.....	32

[1. Introduction](#)

The overall quality, reliability, and performance of VoIP, Real-time Text and Video over IP services rely on the performance and quality of the media path. In order to assure the quality of the delivered media there is a need to monitor the performance of the media transport. One method of monitoring and managing the overall quality of real-time VoIP, Text and Video over IP Services is through monitoring the quality of the media in an active session. This type of "active monitoring" of services is a method of proactively managing the performance and quality of VoIP based services.

The goal of active monitoring is to measure the media quality of a VoIP, Text or Video over IP session. A way to achieve this goal is to request an endpoint to loop media back to the other endpoint and to provide media statistics (e.g., RTCP and RTCP XR information). Another method involves deployment of special endpoints that always loop incoming media back for sessions. Although the latter method has been used and is functional, it does not scale to support large

networks and introduces new network management challenges. Further, it does not offer the granularity of testing a specific endpoint that may be exhibiting problems.

The extension defined in this document introduces new SDP media attributes that enable establishment of media sessions where the media is looped back to the transmitter. The SDP offer/answer model [[RFC3264](#)] is used to establish a loopback connection. Furthermore, this extension provides guidelines on handling RTP [[RFC3550](#)], as well as usage of RTCP [[RFC3550](#)] and RTCP XR [[RFC3611](#)] for reporting media related measurements.

1.1 Use Cases Supported

As a matter of terminology in this document, packets flow from one peer acting as a "loopback source", to the other peer acting as a "loopback mirror", which in turn returns packets to the loopback source. In advance of the session, the peers negotiate to determine which one acts in which role, using the SDP offer/answer exchange. The negotiation also includes details such as the type of loopback to be used.

This specification supports three use cases: "encapsulated packet loopback", "direct loopback", and "media loopback". These are distinguished by the treatment of incoming RTP packets at the loopback mirror.

1.1.1 Encapsulated Packet Loopback

In the encapsulated packet loopback case, the entire incoming RTP packet is encapsulated as payload within an outer RTP packet that is specific to this use case and specified in [Section 7.1](#). The encapsulated packet is returned to the loopback source. The loopback source can generate statistics for one-way path performance up to the RTP level for each direction of travel by examining sequence numbers and timestamps in the encapsulating outer RTP header and the encapsulated RTP packet payload. The loopback source can also play back the returned media content for evaluation.

Because the encapsulating RTP packet header extends the packet size, it could encounter difficulties in an environment where the original RTP packet size is close to the path MTU size. The

encapsulating payload format therefore offers the possibility of RTP-level fragmentation of the returned packets. The use of this facility could affect statistics derived for the return path. In addition, the increased bit rate required in the return direction may affect these statistics more directly in a restricted-bandwidth situation.

1.1.2 Direct Loopback

In the direct loopback case, the loopback mirror copies the payload of the incoming RTP packet into a new RTP packet, using a payload format specific to this use case and specified in [Section 7.2](#). The loopback mirror returns the new packet to the packet source. There is no provision in this case for RTP-level fragmentation.

This use case has the advantage of keeping the packet size the same in both directions. The packet source can compute only two-way path statistics from the direct loopback packet header, but can play back the returned media content.

It has been suggested that the loopback source, knowing that the incoming packet will never be passed to a decoder, can store a timestamp and sequence number inside the payload of the packet it sends to the mirror, then extract that information from the returned direct loopback packet and compute one-way path statistics as in the previous case. Obviously, playout of returned content is no longer possible if this is done.

1.1.3 Media Loopback

In the media loopback case, the loopback mirror submits the incoming packet to a decoder appropriate to the incoming payload type. The packet is taken as close as possible to the analog level, then re-encoded according to an outgoing format determined by SDP negotiation. The reencoded content is returned to the loopback source as an RTP packet with payload type corresponding to the reencoding format.

This usage allows trouble-shooting at the codec level. The capability for path statistics is limited to what is available from RTCP reports.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#).

SDP: Session Description Protocol, as defined in [\[RFC4566\]](#). This document assumes the SDP offer/answer model is followed, per [\[RFC3264\]](#), but does not assume any specific signaling protocol for carrying the SDP.

The following terms are borrowed from [\[RFC3264\]](#) definitions: offer, offerer, answer, answerer, and agent.

3. Overview of Operation

This document defines two loopback 'types', two 'roles', and two encoding formats for loopback. For any given SDP offerer or answerer pair, one side is the source of RTP packets, while the other is the mirror looping packets/media back. Those define the two loopback roles. As the mirror, two 'types' of loopback can be performed: packet-level or media-level. When media-level is used, there is no further choice of encoding format - there is only one format: whatever is indicated for normal media, since the "looping" is performed at the codec level. When packet-level looping is performed, however, the mirror can either send back RTP in an encapsulated format or direct-loopback format. The rest of this document describes these loopback types, roles, and encoding formats, and the SDP offer/answer rules for indicating them.

3.1 SDP Offerer Behavior

An SDP offerer compliant to this memo and attempting to establish a media session with media loopback **MUST** include "loopback" media attributes for each individual media description in the offer message. The offerer **MUST** look for the "loopback" media attributes in the media description(s) of the response from the answer for confirmation that the request is accepted.

3.2 SDP Answerer Behavior

An SDP answerer compliant to this specification and receiving an offer containing media descriptions with the "loopback" media attributes **MUST** acknowledge the request by including the received "loopback" media attributes for each media description in its

answer if it agrees to do the loopback. If the answerer does not want to do loopback or wants to reject the "loopback" request for specific media types, it MAY do so as defined in this section.

An answerer MAY reject an offered stream (either with loopback-source or loopback-mirror) if the loopback-type is not specified, the specified loopback-type is not supported, or the endpoint cannot honor the offer for any other reason. The loopback request MUST be rejected by setting the stream's media port number to zero in the answer as defined in [RFC 3264](#) [[RFC3264](#)], or by rejecting the entire offer (i.e., by rejecting the session request entirely).

Note that an answerer that is not compliant to this specification and which receives an offer with the "loopback" media attributes would ignore the attributes and treat the incoming offer as a normal request. If the offerer does not wish to establish a "normal" RTP session, it would need to terminate the session upon receiving such an answer.

4. New SDP Attributes

Three new SDP media-level attributes are defined: one indicates the type of loopback, and the other two define the role of the agent.

[4.1](#) Loopback Type Attribute

This specification defines a new 'loopback' attribute, which indicates that the agent wishes to perform loopback, and the type of loopback that the agent is able to do. The loopback-type is a property media attribute with the following syntax:

```
a=loopback:<loopback-type>
```

Following is the Augmented BNF [[RFC5234](#)] for loopback-type:

```
Loopback-attr      = "a=loopback:" SP loopback-type
loopback-type      = loopback-choice [1*SP loopback-choice]
loopback-choice    = loopback-type-pkt / loopback-type-media
loopback-type-pkt  = "rtp-pkt-loopback"
loopback-type-media = "rtp-media-loopback"
```

The loopback-type is used to indicate the type of loopback. The loopback-type values are rtp-pkt-loopback, and rtp-media-loopback.

rtp-pkt-loopback: In this mode, the RTP packets are looped back to the sender at a point before the encoder/decoder function in the receive direction to a point after the encoder/decoder function in

the send direction. This effectively re-encapsulates the RTP payload with the RTP/UDP/IP headers appropriate for sending it in the reverse direction. Any type of encoding related functions, such as packet loss concealment, MUST NOT be part of this type of loopback path. In this mode the RTP packets are looped back with a new payload type and format. [Section 7](#) describes the payload formats that MUST be used for this type of loopback. This type of loopback applies to the encapsulated and direct loopback use-cases described in [Section 1.1](#).

rtp-media-loopback: This loopback is activated as close as possible to the analog interface and after the decoder so that the RTP packets are subsequently re-encoded prior to transmission back to the sender. This type of loopback applies to the media loopback use-case described in [Section 1.1.3](#).

[4.2](#) Loopback Role Attributes: loopback-source and loopback-mirror

The loopback role defines two value media attributes that are used to indicate the role of the agent generating the SDP offer or answer. The syntax of the two loopback role media attributes are as follows:

a=loopback-source

and

a=loopback-mirror

loopback-source: This attribute specifies that the entity that generated the SDP is the media source and expects the receiver of the SDP message to act as a loopback-mirror.

loopback-mirror: This attribute specifies that the entity that generated the SDP will mirror (echo) all received media back to the sender of the RTP stream. No media is generated locally by the looping back entity for transmission in the mirrored stream.

The "m=" line in the SDP MUST include all the payload types that will be used during the loopback session. The complete payload space for the session is specified in the "m=" line and the rtpmap attribute is used to map from the payload type number to an encoding name denoting the payload format to be used.

5. Rules for Generating the SDP offer/answer

5.1 Generating the SDP Offer for Loopback Session

If an offerer wishes to make a loopback request, it MUST include both the loopback-type and loopback-role attributes in a valid SDP offer:

```
Example:  m=audio 41352 RTP/AVP 0 8 100
          a=loopback:rtp-media-loopback
          a=loopback-source
          a=rtpmap:0 pcmu/8000
          a=rtpmap:8 pcma/8000
          a=rtpmap:100 G7221/16000/1
```

Since media loopback requires bidirectional RTP, its normal direction mode is "sendrecv"; the "sendrecv" direction attribute MAY be encoded in SDP or not, as per [section 5.1 of \[RFC3264\]](#), since it is implied by default. If either the loopback source or mirror wish to disable loopback use during a session, the direction mode attribute "inactive" MUST be used as per [\[RFC3264\]](#). The direction mode attributes "recvonly" and "sendonly" are incompatible with the loopback mechanism and MUST NOT be indicated when generating an SDP Offer or Answer. When receiving an SDP Offer or Answer, if "recvonly" or "sendonly" is indicated for loopback, the SDP-receiving agent SHOULD treat it as a protocol failure of the loopback negotiation and terminate the session through its normal means (e.g., by sending a SIP BYE if SIP is used).

The offerer may offer more than one loopback-type in the SDP offer. The port number and the address in the offer (m/c= lines) indicate where the offerer would like to receive the media stream(s). The payload type numbers indicate the value of the payload the offerer expects to receive. However, the answer might indicate a subset of payload type numbers from those given in the offer. In that case, the offerer MUST only send the payload types received in the answer, per normal SDP offer/answer rules.

If the offer indicates rtp-pkt-loopback support, the offer MUST also contain either an encapsulated or direct loopback encoding format encoding names, or both, as defined in later sections of this document. If the offer only indicates rtp-media-loopback support, then neither encapsulated nor direct loopback encoding formats apply and they MUST NOT be in the offer.

If loopback-type is rtp-pkt-loopback, the loopback-mirror MUST send and the loopback-source MUST receive the looped back packets

encoded in one of the two payload formats (encapsulated RTP or direct loopback) as defined in [section 7](#).

Example: m=audio 41352 RTP/AVP 0 8 112
 a=loopback:rtp-pkt-loopback
 a=loopback-source
 a=rtpmap:112 encaprtp/8000

Example: m=audio 41352 RTP/AVP 0 8 112
 a=loopback:rtp-pkt-loopback
 a=loopback-source
 a=rtpmap:112 rtploopback/8000

[5.2](#) Generating the SDP Answer for Loopback Session

As with the offer, an SDP answer for loopback MUST follow SDP offer/answer rules for the direction attribute, but directions of "sendonly" or "recvonly" do not apply for loopback operation. \

The port number and the address in the answer (m/c= lines) indicate where the answerer would like to receive the media stream. The payload type numbers indicate the value of the payload types the answerer expects to send and receive.

If an answerer wishes to accept the loopback request it MUST include both the loopback role and loopback type attributes in the answer. When a stream is offered with the loopback-source attribute, the corresponding stream in the response MUST be loopback-mirror and vice versa, provided that answerer is capable of supporting the requested loopback-type.

For example, if the offer contains the loopback-source attribute:

```
m=audio 41352 RTP/AVP 0 8
a=loopback:rtp-media-loopback
a=loopback-source
```

The answer that is capable of supporting the offer MUST contain the loopback-mirror attribute:

```
m=audio 12345 RTP/AVP 0 8
a=loopback:rtp-media-loopback
a=loopback-mirror
```

If a stream is offered with multiple loopback type attributes, the answer MUST include only one of the loopback types that are accepted by the answerer. The answerer SHOULD give preference to the first loopback-type in the SDP offer.

For example, if the offer contains:

```
m=audio 41352 RTP/AVP 0 8 112
a=loopback:rtp-media-loopback rtp-pkt-loopback
a=loopback-source
a=rtpmap:112 encaprtmp/8000
```

The answer that is capable of supporting the offer and chooses to loopback the media using the rtp-media-loopback type MUST contain:

```
m=audio 12345 RTP/AVP 0 8
a=loopback:rtp-media-loopback
a=loopback-mirror
```

As specified in [section 7](#), if the loopback-type is rtp-pkt-loopback, either the encapsulated RTP payload format or direct loopback RTP payload format MUST be used for looped back packets.

For example, if the offer contains:

```
m=audio 41352 RTP/AVP 0 8 112 113
a=loopback:rtp-pkt-loopback
a=loopback-source
a=rtpmap:112 encaprtmp/8000
a=rtpmap:113 rtploopback/8000
```

The answer that is capable of supporting the offer must contain one of the following:

```
m=audio 12345 RTP/AVP 0 8 112
a=loopback:rtp-pkt-loopback
a=loopback-mirror
a=rtpmap:112 encaprtmp/8000
```

```
m=audio 12345 RTP/AVP 0 8 113
a=loopback:rtp-pkt-loopback
a=loopback-mirror
a=rtpmap:113 rtploopback/8000
```

The previous examples used the 'encaprtmp' and 'rtploopback' encoding names, which will be defined in sections [7.1.3](#) and [7.2.3](#).

5.3 Offerer Processing of the SDP Answer

If the received SDP answer does not contain an a=loopback-mirror or a=loopback-source attribute, it is assumed that the loopback

extensions are not supported by the remote agent. This is not a protocol failure, and instead merely completes the SDP offer/answer exchange with whatever normal rules apply; the offerer MAY decide to end the established RTP session (if any) through normal means of the upper-layer signaling protocol (e.g., by sending a SIP BYE).

5.4 Modifying the Session

At any point during the loopback session, either participant MAY issue a new offer to modify the characteristics of the previous session, as defined in [section 8 of RFC 3264](#) [RFC3264]. This also includes transitioning from a normal media processing mode to loopback mode, and vice a versa.

5.5 Establishing Sessions Between Entities Behind NAT

ICE/STUN/TURN provide a general solution to establishing media sessions between entities that are behind NATs, as defined in [RFC5245]. Loopback sessions that involve one or more end points behind NATs SHOULD use these general solutions wherever possible.

Furthermore, if the mirroring entity is behind a NAT, it MUST send some packets to the identified address/port(s) of the peer, in order to open the NAT pinhole. Using ICE this would be accomplished with the STUN connectivity check process, or through a TURN server connection. If ICE is not supported, either [RFC6263] or [Section 10](#) of ICE [RFC5245] SHOULD be followed to open the pinhole and keep the NAT binding alive/refreshed.

Note that for any form of NAT traversal to function, symmetric RTP/RTCP MUST be used. In other words both agents MUST send packets from the same source address and port they receive packets on.

6. RTP Requirements

A loopback source MUST NOT send multiple source streams on the same 5-tuple, since there is no means for the mirror to indicate which is which in its mirrored RTP packets.

A loopback mirror that is compliant to this specification and accepts media with rtp-pkt-loopback loopback-type MUST loopback the incoming RTP packets using either the encapsulated RTP payload format or the direct loopback RTP payload format as defined in [section 7](#) of this specification.

A device that is compliant to this specification and performing the mirroring using the loopback type `rtp-media-loopback` MUST transmit all received media back to the sender, unless congestion feedback or other lower-layer constraints prevent it from doing so. The incoming media MUST be treated as if it were to be played (e.g. the media stream MAY receive treatment from PLC algorithms). The mirroring entity MUST re-generate all the RTP header fields as it would when transmitting media. The mirroring entity MAY choose to encode the loopback media according to any of the media descriptions supported by the offering entity. Furthermore, in cases where the same media type is looped back, the mirroring entity MAY choose to preserve number of frames/packet and bitrate of the encoded media according to the received media.

7. Payload formats for Packet loopback

The payload formats described in this section MUST be used by a loopback-mirror when `'rtp-pkt-loopback'` is the specified loopback-type. Two different formats are specified here - an encapsulated RTP payload format and a direct loopback RTP payload format. The encapsulated RTP payload format should be used when the incoming RTP header information needs to be preserved during the loopback operation. This is useful in cases where loopback source needs to measure performance metrics in both directions. However, this comes at the expense of increased packet size as described in [section 7.1](#). The direct loopback RTP payload format should be used when bandwidth requirement prevents the use of encapsulated RTP payload format.

To keep the implementation of loopback-mirrors simple it is mandated that no payload format other than encapsulated or direct loopback formats can be used in the packets generated by a loopback-mirror. As described in [RFC 3550](#) [[RFC3550](#)], sequence numbers and timestamps in the RTP header are generated with initial random values for security reasons. If this were not mandated and the source payload is sequence number aware, the loopback-mirror will be required to understand that payload format to generate looped back packets that do not violate [RFC 3550](#) [[RFC3550](#)]. Requiring looped back packets to be in one of the two formats means loopback-mirror does not have to look into the actual payload received before generating the loopback packets.

7.1 Encapsulated Payload format

A received RTP packet is encapsulated in the payload section of the RTP packet generated by a loopback-mirror. Each received packet MUST be encapsulated in a separate encapsulating RTP packet; the

encapsulated packet MUST be fragmented only if required (for example: due to MTU limitations).

7.1.1 Usage of RTP Header fields

Payload Type (PT): The assignment of an RTP payload type for this packet format is outside the scope of this document; it is either specified by the RTP profile under which this payload format is used or more likely signaled dynamically out-of-band (e.g., using SDP; [section 7.1.3](#) defines the name binding).

Marker (M) bit: If the received RTP packet is looped back in multiple encapsulating RTP packets, the M bit is set to 1 in every fragment except the last packet, otherwise it is set to 0.

Extension (X) bit: Defined by the RTP Profile used.

Sequence Number: The RTP sequence number SHOULD be generated by the loopback-mirror in the usual manner with a constant random offset as described in [RFC 3550](#) [[RFC3550](#)].

Timestamp: The RTP timestamp denotes the sampling instant for when the loopback-mirror is transmitting this packet to the loopback-source. The RTP timestamp MUST use the same clock rate as that of the encapsulated packet. The initial value of the timestamp SHOULD be random for security reasons (see [Section 5.1 of RFC 3550](#) [[RFC3550](#)]).

SSRC: set as described in [RFC 3550](#) [[RFC3550](#)].

CC and CSRC fields are used as described in [RFC 3550](#) [[RFC3550](#)].

7.1.2 RTP Payload Structure

The outer RTP header of the encapsulating packet MUST be followed by the payload header defined in this section. If the received RTP packet has to be looped back in multiple encapsulating packets due to fragmentation, the encapsulating RTP header in each packet MUST be followed by the payload header defined in this section. The header is devised so that the loopback-source can decode looped back packets in the presence of moderate packet loss [[RFC3550](#)].


```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     receive timestamp                    |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| F | R |  CC  |M|      PT      |      sequence number                  |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|                                     transmit timestamp                    |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|      synchronization source (SSRC) identifier                          |
+==+==+==+==+==+==+==+==+==+==+==+==+==+==+==+==+==+==+==+==+==+==+
|      contributing source (CSRC) identifiers                             |
|                                     ....                                |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

The 12 octets after the receive timestamp are identical to the encapsulated RTP header of the received packet except for the first 2 bits of the first octet. In effect, the received RTP packet is encapsulated by creating a new outer RTP header followed by 4 new bytes of a receive timestamp, followed by the original received RTP header and payload, except that the first two bits of the received RTP header are overwritten as defined here.

Receive Timestamp: 32 bits

The Receive timestamp denotes the sampling instant for when the last octet of the received media packet that is being encapsulated by the loopback-mirror is received from the loopback-source. The Receive timestamp **MUST** be based on the same clock used by the loopback-source. The initial value of the timestamp **SHOULD** be random for security reasons (see [Section 5.1 of RFC 3550](#) [[RFC3550](#)]).

Fragmentation (F): 2 bits

First Fragment (00) /Last Fragment (01) /No Fragmentation(10)/ Intermediate Fragment (11). This field identifies how much of the received packet is encapsulated in this packet by the loopback-mirror. If the received packet is not fragmented, this field is set to 10; otherwise the packet that contains the first fragments sets this field to 00, the packet that contains the last fragment sets this field to 01, all other packets set this field to 11.

7.1.3 Usage of SDP

The payload type number for the encapsulated stream can be negotiated using SDP. There is no static payload type assignment

for the encapsulating stream, so dynamic payload type numbers MUST be used. The binding to the name is indicated by an rtpmap attribute. The name used in this binding is "encaprtp".

The following is an example SDP fragment for encapsulated RTP.

```
m=audio 41352 RTP/AVP 112
a=rtpmap:112 encaprtp/8000
```

7.2 Direct loopback RTP payload format

The direct loopback RTP payload format can be used in scenarios where the 16 byte overhead of the encapsulated payload format is of concern, or simply due to local policy. When using this payload format, the receiver MUST loop back each received RTP packet payload (not header) in a separate RTP packet.

Because a direct loopback format does not retain the original RTP headers, there will be no indication of the original payload-type sent to the mirror, in looped returning packets. Therefore, the loopback source SHOULD only send one payload type per loopback RTP session, if direct mode is used.

7.2.1 Usage of RTP Header fields

Payload Type (PT): The assignment of an RTP payload type for the encapsulating packet format is outside the scope of this document; it is either specified by the RTP profile under which this payload format is used or more likely signaled dynamically out-of-band (e.g., using SDP; [section 7.2.3](#) defines the name binding).

Marker (M) bit: Set to the value in the received packet.

Extension (X) bit: Defined by the RTP Profile used.

Sequence Number: The RTP sequence number SHOULD be generated by the loopback-mirror in the usual manner with a constant random offset.

Timestamp: The RTP timestamp denotes the sampling instant for when the loopback-mirror is transmitting this packet to the loopback-source. The RTP timestamp MUST be based on the same clock as that of the received RTP packet. The initial value of the timestamp SHOULD be random for security reasons (see [Section 5.1 of RFC 3550](#) [[RFC3550](#)]).

SSRC: set as described in [RFC 3550](#) [[RFC3550](#)].

CC and CSRC fields are used as described in [RFC 3550](#) [[RFC3550](#)].

7.2.2 RTP Payload Structure

This payload format does not define any payload specific headers. The loopback-mirror simply copies the RTP payload data from the payload portion of the RTP packet received from the loopback-source.

7.2.3 Usage of SDP

The payload type number for the payload loopback stream can be negotiated using a mechanism like SDP. There is no static payload type assignment for the stream, so dynamic payload type numbers MUST be used. The binding to the name is indicated by an rtpmap attribute. The name used in this binding is "rtploopback".

The following is an example SDP fragment for direct loopback RTP format.

```
m=audio 41352 RTP/AVP 112
a=rtpmap:112 rtploopback/8000
```

8. RTCP Requirements

The use of the loopback attribute is intended for monitoring of media quality of the session. Consequently the media performance information should be exchanged between the offering and the answering entities. An offering or answering agent that is compliant to this specification SHOULD support RTCP per [RFC3550](#) and RTCP-XR per [RFC 3611](#) [[RFC3611](#)]. Furthermore, if the offerer or answerer choose to support RTCP-XR, they SHOULD support RTCP-XR Loss RLE report block, Duplicate RLE report block, Statistics Summary report block, and VoIP Metric Reports Block per sections 4.1, 4.2, 4.6, and 4.7 of [RFC 3611](#) [[RFC3611](#)]. The offerer and the answerer MAY support other RTCP-XR reporting blocks as defined by [RFC 3611](#) [[RFC3611](#)].

9. Congestion Control

All the participants in a loopback session SHOULD implement congestion control mechanisms as defined by the RTP profile under which the loopback mechanism is implemented. For audio video profiles, implementations SHOULD conform to the mechanism defined in [Section 2 of RFC 3551](#).

10. Examples

This section provides examples for media descriptions using SDP for different scenarios. The examples are given for SIP-based transactions and are abbreviated and do not show the complete signaling for convenience.

10.1 Offer for specific media loopback type

An agent sends an SDP offer which looks like:

```
v=0
o=alice 2890844526 2890842807 IN IP4 host.atlanta.example.com
s=-
c=IN IP4 host.atlanta.example.com
t=0 0
m=audio 49170 RTP/AVP 0
a=loopback:rtp-media-loopback
a=loopback-source
a=rtpmap:0 pcmu/8000
```

The agent is offering to source the media and expects the answering agent to mirror the RTP stream per rtp-media-loopback loopback type.

An answering agent sends an SDP answer which looks like:

```
v=0
o=bob 1234567890 1122334455 IN IP4 host.biloxi.example.com
s=-
c=IN IP4 host.biloxi.example.com
t=0 0
m=audio 49270 RTP/AVP 0
a=loopback:rtp-media-loopback
a=loopback-mirror
a=rtpmap:0 pcmu/8000
```

The answerer is accepting to mirror the media from the offerer at the media level.

10.2 Offer for choice of media loopback type

An agent sends an SDP offer which looks like:

```
v=0
```



```
o=alice 2890844526 2890842807 IN IP4 host.atlanta.example.com
s=-
c=IN IP4 host.atlanta.example.com
t=0 0
m=audio 49170 RTP/AVP 0 112 113
a=loopback:rtp-media-loopback rtp-pkt-loopback
a=loopback-source
a=rtpmap:0 pcmu/8000
a=rtpmap:112 encaprtp/8000
a=rtpmap:113 rtploopback/8000
```

The offerer is offering to source the media and expects the answerer to mirror the RTP stream at either the media or rtp level.

An answering agent sends an SDP answer which looks like:

```
v=0
o=box 1234567890 1122334455 IN IP4 host.biloxi.example.com
s=-
c=IN IP4 host.biloxi.example.com
t=0 0
m=audio 49270 RTP/AVP 0 112
a=loopback:rtp-pkt-loopback
a=loopback-mirror
a=rtpmap:0 pcmu/8000
a=rtpmap:112 encaprtp/8000
```

The answerer is accepting to mirror the media from the offerer at the packet level using the encapsulated RTP payload format.

[10.3](#) Answerer rejecting loopback media

An agent sends an SDP offer which looks like:

```
v=0
o=alice 2890844526 2890842807 IN IP4 host.atlanta.example.com
s=-
c=IN IP4 host.atlanta.example.com
t=0 0
m=audio 49170 RTP/AVP 0
a=loopback:rtp-media-loopback
a=loopback-source
a=rtpmap:0 pcmu/8000
```

The offerer is offering to source the media and expects the answerer to mirror the RTP stream at the media level.

An answering agent sends an SDP answer which looks like:

```
v=0
o=bob 1234567890 1122334455 IN IP4 host.biloxi.example.com
s=-
c=IN IP4 host.biloxi.example.com
t=0 0
m=audio 0 RTP/AVP 0
a=rtpmap:0 pcmu/8000
```

Note in this case the answerer did not indicate loopback support, although it could have and still used a port number of 0 to indicate it does not wish to accept that media session.

Alternatively, the answering agent could have simply rejected the entire SDP offer through some higher-layer signaling protocol means (e.g., by rejecting the SIP INVITE request if the SDP offer was in the INVITE).

11. Security Considerations

The security considerations of [[RFC3264](#)] and [[RFC3550](#)] apply.

Given that media loopback may be automated without the end user's knowledge, the answerer of the media loopback should be aware of denial of service attacks. It is recommended that session requests for media loopback be authenticated and the frequency of such sessions limited by the answerer.

If the higher-layer signaling protocol were not authenticated, a malicious attacker could create a session between two parties the attacker wishes to target, with each party acting as the loopback-mirror to the other, of rtp-pkt-loopback type. A few RTP packets sent to either party would then infinitely loop among the two, as fast as they could process them, consuming their resources and network bandwidth.

Furthermore, media-loopback provides a means of attack indirection, whereby a malicious attacker creates a loopback session as the loopback-source, and uses the mirror to reflect the attacker's packets against a target - perhaps a target the attacker could not reach directly, such as one behind a firewall for example. Or the attacker could initiate the session as the loopback-mirror, in the hopes of making the peer generate media against another target.

If end-user devices such as mobile phones answer loopback requests without authentication and without notifying the end-user, then an attacker could cause the battery to drain, and possibly deny the end-user normal phone service or cause network data usage fees.

This could even occur naturally if a legitimate loopback session does not terminate properly and the end device does not have a timeout mechanism for such.

For the reasons noted above, end user devices SHOULD provide a means of indicating to the human user that the device is in a loopback session, even if it is an authenticated session. Devices which answer or generate loopback sessions SHOULD either perform keepalive/refresh tests of the session state through some means, or time out the session automatically.

12. Implementation Considerations

The media loopback approach described in this document is a complete solution that would work under all scenarios. However, it is believed that the solution may not be light-weight enough for the common case. In light of this concern, this section clarifies which features of the loopback proposal MUST be implemented for all implementations and which features MAY be deferred if the complete solution is not desired.

All implementations MUST at least support the rtp-pkt-loopback mode for loopback-type, with direct media loopback payload encoding. In addition, for the loopback role, all implementations of an SDP offerer MUST at least be able to act as a loopback-source.

13. IANA Considerations

13.1 SDP Attributes

This document defines three new media-level SDP attributes. IANA has registered the following attributes:

Contact name:	Kaynam Hedayat
<kaynam.hedayat@exfo.com>.	
Attribute name:	"loopback".
Type of attribute:	Media level.
Subject to charset:	No.
Purpose of attribute:	The 'loopback' attribute is used to indicate the type of media loopback.
Allowed attribute values:	The parameters to 'loopback' may be one or more of "rtp-pkt-loopback" and "rtp-media-loopback". See section 5 of this document for syntax.

Contact name: Kaynam Hedayat
<kaynam.hedayat@exfo.com>.
Attribute name: "loopback-source".
Type of attribute: Media level.
Subject to charset: No.
Purpose of attribute: The 'loopback-source' attribute specifies that the sender is the media source and expects the receiver to act as a loopback-mirror.
Allowed attribute values: None.

Contact name: Kaynam Hedayat
<kaynam.hedayat@exfo.com>.
Attribute name: "loopback-mirror".
Type of attribute: Media level.
Subject to charset: No.
Purpose of attribute: The 'loopback-mirror' attribute specifies that the receiver will mirror (echo) all received media back to the sender of the RTP stream.
Allowed attribute values: None.

13.2 MIME Types

The IANA has registered the following MIME types:

13.2.1 audio/encaprtmp

To: ietf-types@iana.org

Subject: Registration of media type audio/encaprtmp

Type name: audio

Subtype name: encaprtmp

Required parameters:

rate: RTP timestamp clock rate, which is equal to the sampling rate. The typical rate is 8000; other rates may be specified.

Optional parameters: none

Encoding considerations: This media type is framed binary data.

Security considerations: See [Section 12](#) of this document.

Interoperability considerations: none

Published specification: This MIME type is described fully within this document.

Applications which use this media type: Applications wishing to monitor and ensure the quality of transport to the edge of a given VoIP, Real-Time Text or Video Over IP Service.

Additional information: none

Person & email address to contact for further information:

Kaynam Hedayat
EMail: kaynam.hedayat@exfo.com

Intended usage: COMMON

Restrictions on usage: This media type depends on RTP framing, and hence is only defined for transfer via RTP. Transfer within other framing protocols is not defined at this time.

Author:
Kaynam Hedayat.

Change controller: IETF Audio/Video Transport working group delegated from the IESG.

[13.2.2](#) **video/encaprtmp**

To: ietf-types@iana.org

Subject: Registration of media type video/encaprtmp

Type name: video

Subtype name: encaprtmp

Required parameters:

rate: RTP timestamp clock rate, which is equal to the sampling rate. The typical rate is 8000; other rates may be specified.

Optional parameters: none

Encoding considerations: This media type is framed binary data.

Security considerations: See [Section 12](#) of this document.

Interoperability considerations: none

Published specification: This MIME type is described fully within this document.

Applications which use this media type: Applications wishing to monitor and ensure the quality of transport to the edge of a given VoIP, Real-Time Text or Video Over IP Service.

Additional information: none

Person & email address to contact for further information:

Kaynam Hedayat
EMail: kaynam.hedayat@exfo.com

Intended usage: COMMON

Restrictions on usage: This media type depends on RTP framing, and hence is only defined for transfer via RTP. Transfer within other framing protocols is not defined at this time.

Author:
Kaynam Hedayat.

Change controller: IETF Audio/Video Transport working group delegated from the IESG.

[13.2.3](#) **text/encaprtmp**

To: ietf-types@iana.org

Subject: Registration of media type text/encaprtmp

Type name: text

Subtype name: encaprtmp

Required parameters:

rate:RTP timestamp clock rate, which is equal to the sampling rate. The typical rate is 8000; other rates may be specified.

Optional parameters: none

Encoding considerations: This media type is framed binary data.

Security considerations: See [Section 12](#) of this document.

Interoperability considerations: none

Published specification: This MIME type is described fully within this document.

Applications which use this media type: Applications wishing to monitor and ensure the quality of transport to the edge of a given VoIP, Real-Time Text or Video Over IP Service.

Additional information: none

Person & email address to contact for further information:

Kaynam Hedayat
EMail: kaynam.hedayat@exfo.com

Intended usage: COMMON

Restrictions on usage: This media type depends on RTP framing, and hence is only defined for transfer via RTP. Transfer within other framing protocols is not defined at this time.

Author:

Kaynam Hedayat.

Change controller: IETF Audio/Video Transport working group delegated from the IESG.

[13.2.4](#) **application/encaprtsp**

To: ietf-types@iana.org

Subject: Registration of media type
application/encaprtsp

Type name: application

Subtype name: encaprtsp

Required parameters:

rate:RTP timestamp clock rate, which is equal to the
sampling rate. The typical rate is 8000; other rates
may be specified.

Optional parameters: none

Encoding considerations: This media type is framed
binary data.

Security considerations: See [Section 12](#) of this document.

Interoperability considerations: none

Published specification: This MIME type is described fully
within this document.

Applications which use this media type: Applications wishing
to monitor and ensure the quality of transport to the
edge of a given VoIP, Real-Time Text or Video Over IP
Service.

Additional information: none

Person & email address to contact for further information:

Kaynam Hedayat
EMail: kaynam.hedayat@exfo.com

Intended usage: COMMON

Restrictions on usage: This media type depends on RTP
framing, and hence is only defined for transfer via
RTP. Transfer within other framing protocols is not
defined at this time.

Author:

Kaynam Hedayat.

Change controller: IETF Audio/Video Transport working
group delegated from the IESG.

13.2.5 **audio/rtploopback**

To: ietf-types@iana.org

Subject: Registration of media type audio/rtploopback

Type name: audio

Subtype name: rtploopback

Required parameters:

 rate:RTP timestamp clock rate, which is equal to the
 sampling rate. The typical rate is 8000; other rates
 may be specified.

Optional parameters: none

Encoding considerations: This media type is framed
 binary data.

Security considerations: See [Section 12](#) of this document.

Interoperability considerations: none

Published specification: This MIME type is described fully
 within this document.

Applications which use this media type: Applications wishing
 to monitor and ensure the quality of transport to the
 edge of a given VoIP, Real-Time Text or Video Over IP
 Service.

Additional information: none

Person & email address to contact for further information:

 Kaynam Hedayat
 EMail: kaynam.hedayat@exfo.com

Intended usage: COMMON

Restrictions on usage: This media type depends on RTP
 framing, and hence is only defined for transfer via
 RTP. Transfer within other framing protocols is not
 defined at this time.

Author:

Kaynam Hedayat.

Change controller: IETF Audio/Video Transport working group delegated from the IESG.

13.2.6 video/rtploopback

To: ietf-types@iana.org

Subject: Registration of media type video/rtploopback

Type name: video

Subtype name: rtploopback

Required parameters:

rate:RTP timestamp clock rate, which is equal to the sampling rate. The typical rate is 8000; other rates may be specified.

Optional parameters: none

Encoding considerations: This media type is framed binary data.

Security considerations: See [Section 12](#) of this document.

Interoperability considerations: none

Published specification: This MIME type is described fully within this document.

Applications which use this media type: Applications wishing to monitor and ensure the quality of transport to the edge of a given VoIP, Real-Time Text or Video Over IP Service.

Additional information: none

Person & email address to contact for further information:

Kaynam Hedayat
EMail: kaynam.hedayat@exfo.com

Intended usage: COMMON

Restrictions on usage: This media type depends on RTP framing, and hence is only defined for transfer via

RTP. Transfer within other framing protocols is not defined at this time.

Author:

Kaynam Hedayat.

Change controller: IETF Audio/Video Transport working group delegated from the IESG.

[13.2.7](#) **text/rtploopback**

To: ietf-types@iana.org

Subject: Registration of media type text/rtploopback

Type name: text

Subtype name: rtploopback

Required parameters:

rate: RTP timestamp clock rate, which is equal to the sampling rate. The typical rate is 8000; other rates may be specified.

Optional parameters: none

Encoding considerations: This media type is framed binary data.

Security considerations: See [Section 12](#) of this document.

Interoperability considerations: none

Published specification: This MIME type is described fully within this document.

Applications which use this media type: Applications wishing to monitor and ensure the quality of transport to the edge of a given VoIP, Real-Time Text or Video Over IP Service.

Additional information: none

Person & email address to contact for further information:

Kaynam Hedayat

EMail: kaynam.hedayat@exfo.com

Intended usage: COMMON

Restrictions on usage: This media type depends on RTP framing, and hence is only defined for transfer via RTP. Transfer within other framing protocols is not defined at this time.

Author:
Kaynam Hedayat.

Change controller: IETF Audio/Video Transport working group delegated from the IESG.

13.2.8 application/rtploopback

To: ietf-types@iana.org

Subject: Registration of media type
application/rtploopback

Type name: application

Subtype name: rtploopback

Required parameters:

rate: RTP timestamp clock rate, which is equal to the sampling rate. The typical rate is 8000; other rates may be specified.

Optional parameters: none

Encoding considerations: This media type is framed binary data.

Security considerations: See [Section 12](#) of this document.

Interoperability considerations: none

Published specification: This MIME type is described fully within this document.

Applications which use this media type: Applications wishing to monitor and ensure the quality of transport to the edge of a given VoIP, Real-Time Text or Video Over IP Service.

Additional information: none

Person & email address to contact for further information:

Kaynam Hedayat
EMail: kaynam.hedayat@exfo.com

Intended usage: COMMON

Restrictions on usage: This media type depends on RTP framing, and hence is only defined for transfer via RTP. Transfer within other framing protocols is not defined at this time.

Author:
Kaynam Hedayat.

Change controller: IETF Audio/Video Transport working group delegated from the IESG.

14. Acknowledgements

This document's editor would like to thank the original authors of the document: Kaynam Hedayat, et al. The editor has made fairly insignificant changes in the end. Also, we'd like to thank Magnus Westerlund, Miguel Garcia, Flemming Andreason, Gunnar Hellstrom, Emil Ivov and Dan Wing for their feedback, comments and suggestions.

15. Normative References

- [RFC3264] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with the Session Description Protocol (SDP)", [RFC 3264](#), June 2002.
- [RFC3550] Schulzrinne, H., Casner, S., Frederick, R. and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, [RFC 3550](#), July 2003.
- [RFC3611] Almeroth, K., Caceres, R., Clark, A., Cole, R., Duffield, N., Friedman, T., Hedayat, K., Sarac, K. and M. Westerlund, "RTP Control Protocol Extended Reports (RTCP XR)", [RFC 3611](#), November 2003.

- [RFC5234] Crocker, P. Overell, "Augmented ABNF for Syntax Specification: ABNF", [RFC 5234](#), October 2005.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2736] Handley, M., Perkins, C., "Guidelines for Writers of RTP Payload Format Specifications", [RFC 2736](#), [BCP 0036](#), December 1999.
- [RFC3551] Schulzrinne, H., Casner, S., "RTP Profile for Audio and Video Conferences with Minimal Control", STD 65, [RFC 3551](#), July 2003.
- [RFC4566] Handley, M., Jacobson, V., Perkins, C., "SDP: Session Description Protocol", [RFC 4566](#), July 2006.
- [RFC4855] Casner, S., "Media Type Registration of RTP Payload Formats", [RFC 4855](#), February 2007.

[16](#). Informative References

- [RFC5245] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", [RFC 5245](#), April 2010.
- [RFC6263] Marjou, X., Sollaud, A., "Application Mechanism for Keeping Alive the NAT Mappings Associated with RTP / RTP Control Protocol (RTCP) Flows", [RFC 6263](#), June 2011.

Authors' Addresses

Hadriel Kaplan
Acme Packet
100 Crosby Drive
Bedford, MA 01730
USA

EMail: hkaplan@acmepacket.com
URI: <http://www.acmepacket.com>

Kaynam Hedayat
EXFO
285 Mill Road
Chelmsford, MA 01824
US

EMail: kaynam.hedayat@exfo.com
URI: <http://www.exfo.com/>

Nagarjuna Venna
Saperix
738 Main Street, #398
Waltham, MA 02451
US

EMail: vnagarjuna@saperix.com
URI: <http://www.saperix.com/>

Paul E. Jones
Cisco Systems, Inc.
7025 Kit Creek Rd.
Research Triangle Park, NC 27709
US

EMail: paulej@packetizer.com
URI: <http://www.cisco.com/>

Arjun Roychowdhury
Hughes Systique Corp.
15245 Shady Grove Rd, Ste 330
Rockville MD 20850
US

EMail: arjun@hsc.com
URI: <http://www.hsc.com/>

Chelliah SivaChelvan
Cisco Systems, Inc.
2200 East President George Bush Turnpike
Richardson, TX 75082
US

EMail: chelliah@cisco.com
URI: <http://www.cisco.com/>

Nathan Stratton
BlinkMind, Inc.
2027 Briarchester Dr.
Katy, TX 77450

E-Mail: nathan@robotics.net

URI: <http://www.robotics.net/>